An approach based on IoT and machine learning for monitoring patients on healthcare centers

Thiago Bezerra[1,2], Marcos Vinicius[2], Aline Ciane[2], Gustavo Callou[3], Cleunio França[1], Eduardo Tavares[1]

[1] Centro de Informática,
Universidade Federal de Pernambuco, Recife, Brasil

[2]Instituto Federal de Pernambuco, Palmares, Brasil

[3]Departamento de Computação,
Universidade Federal Rural de Pernambuco, Recife, Brasil

{tvb,cbff,eagt}@cin.ufpe.br, {acgs8,mvvm@}discente.ifpe.edu.br

gustavo.callou@ufrpe.br

Abstract. *Internet of Things (IoT) enables remote monitoring using existing networks, integrating diverse computing systems and communication protocols. This adaptability makes IoT applicable across domains, including healthcare. Hospitals leverage the Internet of Medical Things (IoMT) to collect and transmit patient data, enabling real-time diagnoses, operational efficiency, and personalized care. This paper presents a patient monitoring system for healthcare centers using IoMT and machine learning (ML). The system is capable of performing facial recognition of patients, identifying movement patterns associated with falls, using cameras, and sending alerts to caregivers or healthcare professionals. Results demonstrate the system's feasibility in enhancing healthcare outcomes. The system achieved a precision of 95% for facial recognition and 97% for fall detection.*

1. Introduction

The rapid expansion of the Internet of Things (IoT) is transforming the interconnection between physical and virtual devices, leveraging a diverse set of technologies. This evolution has profound implications across multiple domains, including healthcare, industry, and agriculture [Verma et al. 2022]. Notably, applications in resource-constrained environments demonstrate significant potential. The growing adoption of IoT-based medical technologies has driven substantial investments, reflecting an increasing demand for innovation in healthcare and remote patient monitoring [Cisco and Internet 2020]. This trend has led to the emergence of the concept of the Internet of Medical Things (IoMT). In healthcare settings, hospitals are integrating IoMT systems to enhance data collection and transmission, thereby improving the continuous monitoring of critically ill patients [Verma et al. 2022].

The convergence of IoT and ML has introduced advanced monitoring capabilities, particularly for elderly individuals. These technologies enable intelligent and proactive health surveillance by identifying and responding to conditions that may pose health risks [Bharadwaj et al. 2021]. Cameras equipped with ML algorithms can analyze movement patterns and behaviors, distinguishing between normal activities and abnormal events such as falls. This capability allows for the immediate generation of alerts, facilitating timely intervention by caregivers or healthcare systems [Bharadwaj et al. 2021].

Beyond fall detection, facial recognition in IoMT systems expands the potential applications in healthcare. Continuous patient identification and tracking ensure that medical treatments and monitoring are accurately associated with the correct individuals. Facial recognition can also contribute to emotional and cognitive assessments by analyzing facial expressions to detect signs of pain, discomfort, or distress, which is particularly relevant in elderly care and mental health applications. Moreover, integrating IoT and ML enables real-time, continuous analysis of behavioral and health-related patterns. ML algorithms, trained on patient movement data, can detect early indicators of health deterioration, such as irregular movement patterns or prolonged inactivity. As a result, IoMT systems can autonomously trigger alerts, enabling caregivers to intervene before conditions escalate.

This work proposes a patient monitoring system for healthcare environments, integrating IoT and ML to perform facial recognition and detect movement patterns associated with falls, generating real-time alerts for caregivers and healthcare professionals. The primary objective of the research is to develop a system that detects falls while also identifying and tracking patients in real-time, improving patient safety and care quality. The proposal stands out from existing literature, which primarily focuses on isolated systems for fall detection or facial recognition, by integrating these functionalities in a coordinated manner, ensuring a rapid and precise response in critical situations. Additionally, integration with hospital systems enables automatic alerts to nurses and doctors, ensuring a prompt response to critical situations, an area where more effective solutions are still lacking in current literature.

The remainder of this paper is organized as follows. Section 2 presents the related works available in the literature. Section 3 provides an overview of prominent concepts for understanding the proposed approach. In Section 4, the methodology is explained. Section 5 presents the IoMT architectures under analysis, and Section 6 demonstrates the practical application of the proposed solution e presents the results of the model evaluation. Finally, Section 7 concludes this study and shows future research directions.

2. Related work

This section presents a review of recent research efforts that apply IoMT, ML, and deep learning (DL) techniques to patient monitoring in healthcare contexts. The studies are organized around key functionalities such as physiological signal analysis, facial recognition for identification, and fall detection. These works contribute to the understanding of current capabilities and limitations in the field and serve as a basis for identifying the distinctive contributions of the approach proposed in this study. Rajan and Nadar [Rajan Jeyaraj and Nadar 2022] propose an IoMT system for monitoring physiological data in elderly patients. Their approach incorporates a ML algorithm for feature extraction and anomaly detection, achieving an average accuracy of 97.5% for signals such as Electroencephalogram (EEG), Electrocardiogram (ECG), temperature, and pulse. Additionally, the system provides recommendations for patients to seek medical assistance.

Masud *et al.* [Masud et al. 2020] introduce a tree-based deep learning model for facial recognition in IoMT applications, integrating cloud computing to enhance scalability. The model achieved up to 99% accuracy on public datasets while maintaining a reduced computational cost. Evaluations on datasets such as Facial Expression Interface (FEI), Olivetti Research Laboratory (ORL), and Labeled Faces in the Wild (LFW) demonstrated superior efficiency compared to conventional models. Similarly, Bisogni *et al.* [Bisogni et al. 2022] present a facial expression recognition system (FERS) aimed at improving healthcare services by detecting and classifying patient emotions in real time. Their approach utilizes multi-resolution facial images processed through convolutional neural network (CNN) architectures, surpassing the performance of existing methods in public dataset evaluations.

Altameem *et al.* [Altameem and Altameem 2020] propose a multimodal visualization analysis (MMVA) technique for patient monitoring, focusing on facial expression identification through a three-layer CNN model. This method achieves a recognition accuracy of 95.702%, demonstrating its effectiveness in healthcare applications. Additionally, Talaat [Talaat 2023] presents a deep convolutional neural network (DCNN)-based system designed for facial recognition and emotion detection in autistic children, facilitating early autism diagnosis. By leveraging fog computing and IoT, the system minimizes latency and enhances real-time detection accuracy. Despite using a relatively small dataset, the approach attained an accuracy of 99.99%.

Doan [Doan 2022] introduces an intelligent remote monitoring system that detects and alerts healthcare providers about abnormal patient actions, such as falls. The system employs the MediaPipe Pose library for human body pose estimation, integrated with an LSTM (Long Short-Term Memory) network for patient action classification. Experimental results indicate that the system achieved a classification accuracy of 96.84%, offering advantages such as cost-effectiveness, ease of deployment, and applicability in remote patient monitoring. Koçak [Koçak and Çetin 2021] proposed a model for real-time fall detection in hospitalized patients, using deep learning techniques based on LSTM (Long Short-Term Memory) applied to IoT sensor data. The model achieved a 98% F1 score in real-time fall detection. Additionally, a mobile application was successfully developed to notify caregivers about patient fall events.

Table 1 presents a comparative summary of related works found in the literature on IoT and machine learning for monitoring patients in healthcare centers. A comparative analysis of the strategies adopted in each paper and the key metrics of interest can be conducted through this table. Additionally, it highlights work that performs fall detection and facial recognition. Some studies focused only on fall detection [Doan 2022], [Koçak and Çetin 2021], while only a few evaluate facial recognition [Masud et al. 2020], [Bisogni et al. 2022], [Altameem and Altameem 2020], [Talaat 2023] of monitoring patients on healthcare centers.

Unlike previous studies, our approach extends beyond facial recognition and fall detection by incorporating a patient monitoring system. While prior studies focus on specific aspects such as physiological monitoring, facial recognition, or action classification, our work combines pose detection, fall identification, and movement pattern analysis to assess patient risks. Additionally, it enables real-time alerts to caregivers and healthcare professionals, ensuring immediate attention to potential incidents.

Table 1. Summary of related work.

Work	Facial Recognition	Fall Detection	Pose Detection	Real-time Alerts	ML/DL Techniques
Rajan and Nadar	-	-	-	-	ML
Masud et al.	√	-	-	-	Tree-based DL
Bisogni et al.	√	-	-	-	CNN
Altameem et al.	√	-	-	-	CNN
Talaat	√	-	-	-	DCNN
Doan	-	√	√	√	LSTM
Koçak	-	√	-	√	LSTM
This Work	√	√	√	√	ML

3. Background

This section introduces essential concepts to better understand this work.

3.1. IoT architecture

IoT allows remote monitoring of environmental elements using existing network infrastructures, making it easier to connect distinct computer systems. Such an integration allows efficient data gathering, monitoring, and processing, which are essential for many applications, such as smart homes, healthcare and agriculture. IoT also facilitates dynamic and autonomous communication between devices, which is an important feature for operation in harsh locations [Verma et al. 2022].

A typical IoT architecture (Figure 1) is generally structured into four distinct layers: device, communication, processing, and presentation. is responsible for data gathering and includes sensors, microcontrollers, and other hardware components that may monitor an environment, such as temperature, humidity, motion, or health metrics. These devices are the core of an IoT network, providing the essential raw data required for further analysis and decision-making. The communication layer then acts as the channel in which data is transmitted from the devices to the next processing stages. This layer utilizes standard communication protocols to ensure secure and efficient data transfer to other systems [Verma et al. 2022].

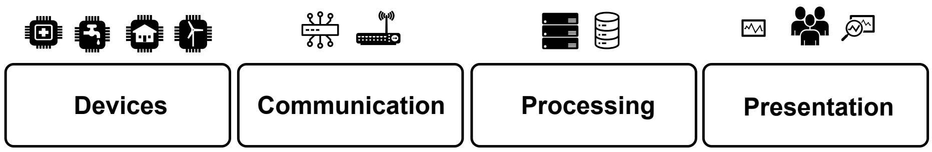

Figure 1. IoT Basic Architecture

The processing layer is responsible for managing, analyzing, and storing the collected data. The layer may perform data aggregation, filtering, and computation, often using cloud or edge computing. Processing layer transforms raw data into valuable insights, which can then be adopted to provide system services or trigger specific actions.

In addition, the processing layer can incorporate ML algorithms, data analytics, and artificial intelligence to improve decision-making and automate responses to specific events or conditions. Finally, the presentation layer provides an interface for users to interact with the IoT system. The layer includes dashboards, visualization tools, mobile applications, and web interfaces that present processed data.

3.2. Machine Learning

Machine learning (ML) is a specialized branch of artificial intelligence (AI) that focuses on creating algorithms that are capable of learning from data, recognizing patterns, and making decisions with minimal human intervention. Unlike traditional programming, where explicit instructions are provided for every scenario, ML systems improve their performance by analyzing past experiences and dynamically adapting to new information. One of the core strengths of ML lies in its ability to handle vast and complex datasets. Whether used for predictive analytics in finance, fraud detection in banking, or recommendation systems in e-commerce, ML algorithms continuously refine their accuracy over time. In healthcare, ML models assist in diagnosing diseases by analyzing patient records and medical images, while in marketing, they help businesses personalize content based on user preferences [Mitchell and Mitchell 1997].

ML techniques include supervised learning, where models are trained on labeled datasets; unsupervised learning, which identifies patterns in unlabeled data; and reinforcement learning, where algorithms learn by interacting with their environment and receiving feedback. The versatility and adaptability of ML make it an indispensable tool in today's digital age, with applications ranging from natural language processing and robotics to autonomous systems and cybersecurity. As research in ML continues to progress, newer methodologies such as deep learning and neural networks are pushing the boundaries of what machines can achieve. With an increasing amount of data being generated every day, ML will play an even more vital role in optimizing processes, enhancing decision-making, and driving technological breakthroughs across industries [Mitchell and Mitchell 1997].

3.3. Confusion Matrix

The confusion matrix (see Table 2) is a widely used tool for evaluating classification models in machine learning and statistics. This tool provides a tabular representation of the predictions of a model compared to actual values, enabling performance analysis. Organized into four main categories - true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN), the matrix helps to understand metrics such as precision, recall, and F1 score [Visa et al. 2011]. The confusion matrix for this classification model is shown below.

Table 2. Confusion Matrix

	Predictions Positive	**Predictions Negative**	**Total**
Positive	True Positive (TP)	False Negative (FN)	Total Positives
Negative	False Positive (FP)	True Negative (TN)	Total Negatives
Total	Pred. Positives	Pred. Negatives	Overall Total

A true positive (TP) occurs when the model correctly predicts a positive instance, meaning that both the prediction and the actual value are positive. Conversely, a true negative (TN) represents a correctly predicted negative instance, where both the prediction and the actual value are negative. A false positive (FP), also known as a Type I error, happens when the model incorrectly predicts a positive instance when the actual value is negative. Finally, a false negative (FN), or Type II error, occurs when the model incorrectly predicts a negative instance when the actual value is actually positive. These four categories form the basis for evaluating the effectiveness of classification models.

3.4. Classification metrics

Classification metrics are widely used to evaluate the performance of machine learning models in classification tasks. Each metric highlights a specific aspect of the model's behavior. Precision (Equation 1) measures the proportion of correctly predicted positive instances among all instances classified as positive. Recall, also known as Sensitivity (Equation 2), quantifies the model's ability to correctly identify all actual positive cases. The F1-score (Equation 3) represents the harmonic mean of precision and recall, offering a single measure that balances both metrics, especially relevant in imbalanced datasets. Accuracy (Equation 4) evaluates the overall correctness of the model by considering both true positives and true negatives among all predictions [Erickson and Kitamura 2021].

$$\text{Precision} = \frac{TP}{TP + FP} \tag{1}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{2}$$

$$F1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{3}$$

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{4}$$

4. Methodology

This section presents the methodology for implementing an intelligent patient monitoring system integrating IoT and ML for fall detection and facial recognition in healthcare centers (see Figure 2). The approach follows sequential steps to ensure system application. The first step, system understanding, involves analyzing the system's components, interactions, and technical requirements. Python is used as the programming language, with the OpenCV library and cvzone module for pose estimation and facial recognition. A Raspberry Pi 5 with an integrated camera is used for real-time image capture and processing.

In the data collection step, facial images are acquired under various conditions to improve model robustness. The images are preprocessed for consistency, optimizing facial recognition accuracy. The model training step employs the Local Binary Patterns Histograms (LBPH) algorithm for feature extraction, encoding facial patterns, and associating them with identity labels for precise recognition.

The fall detection implementation and facial recognition step involves designing an algorithm that continuously monitors a patient's posture using pose estimation techniques with facial recognition. The system analyzes key body points, particularly the vertical positions of the head and knees, to identify potential falls. When the detected posture deviates significantly from an upright position, crossing a predefined threshold indicative of a fall, the system triggers an alert.

The system evaluation step verifies the accuracy of the facial recognition and fall detection processes. The system's performance is assessed to ensure that facial recognition is accurate and that alerts are triggered correctly when a fall is detected. If the results are inconclusive or the system fails to meet expectations, the data collection step is revisited to identify potential flaws or areas for improvement. Finally, the last step of the methodology corresponds to the analysis of results.

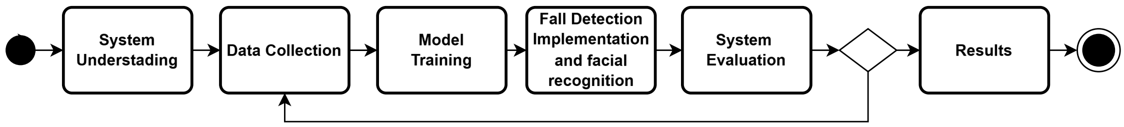

Figure 2. Methodology

5. IoMT Architecture

IoMT (Internet of Medical Things) architectures refer to smart environments equipped with electronic devices and sensors designed to monitor patients' physiological signals. These architectures provide physicians with real-time data access and analysis, thereby enhancing decision-making and improving patient care [Hireche et al. 2022].

Typically, such systems encompass a variety of components, including communication protocols, data storage, data analytics, visualization tools, and both hardware and software elements, all of which enable medical personnel to monitor and remotely manage patient health. The architecture can gather data from a diverse array of sources, such as wearable and mobile devices, offering a comprehensive view of a patient's health status [Askar et al. 2022]. Figure 3 illustrates the IoMT architecture adopted in this work, which, based on [Vishnu et al. 2020], is composed of four integrated layers, specifically designed for facial recognition and fall detection.

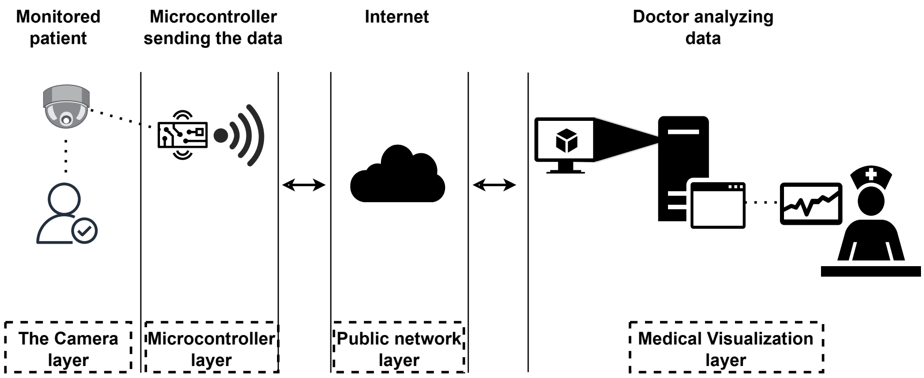

Figure 3. IoMT Architecture

The camera layer is responsible for capturing the initial data. It records images that are used for both facial recognition and fall detection, achieved through the analysis of body movements and positions. This layer acts as the primary interface with the patient, ensuring the accurate collection of relevant information. The collected data is then transmitted to the microcontroller layer, which plays a pivotal role in executing ML algorithms. Once processed, the data is transferred via the public network layer (i.e., the Internet), which serves as the system's core for connectivity and storage. Finally, the medical visualization layer connects healthcare professionals to the generated data. Through interactive dashboards, doctors and caregivers can access and effectively analyze the information.

6. Case study

This section presents a case study illustrating the applicability of the proposed IoMT architecture. The proposed solution provides a patient monitoring system that integrates IoT technology and ML to demonstrate the operation of a fall detection system. Figure 4 presents two distinct aspects of the process used to identify patient falls. In Figure 4(a), the facial recognition stage is illustrated, where the patient is registered to enable continuous identification and activity monitoring.

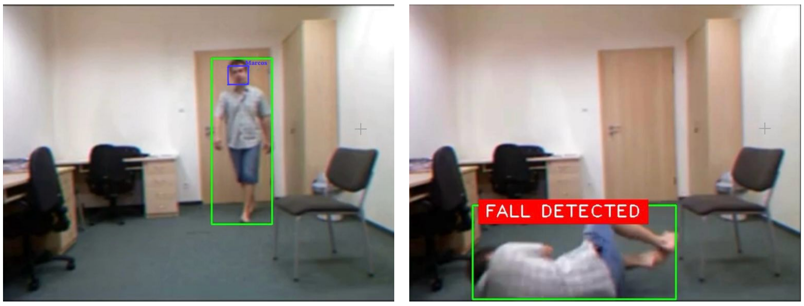

a) Facial recognition and positioning **b) Fall detection and alert generation**

Figure 4. System operation

ML algorithms ensure accurate facial recognition under varying lighting conditions and viewing angles. Additionally, the system tracks the patient's position within the monitored environment, facilitating adequate supervision. By leveraging computer vision and image analysis algorithms, patient movements are continuously monitored, allowing the system to detect spatial positioning and behavioral patterns that may indicate a heightened risk of falls. Figure 4(b) illustrates the moment when a fall is detected. In addition to facial recognition, the system integrates movement pattern recognition algorithms capable of identifying fall-related behaviors. Upon detection, an alert is automatically triggered, notifying caregivers or healthcare professionals in real-time. This individualized identification process ensures that each incident is logged separately, improving reliability and accountability in critical scenarios.

6.1. Model evaluation

This subsection evaluates the effectiveness of the model, focusing on facial recognition and fall detection. The assessment utilizes well-established classification metrics derived from the confusion matrix to quantify the model's ability to differentiate between predefined classes. For fall detection, the model must accurately classify instances as "No Fall" (Class 0) or "Fall" (Class 1), while for facial recognition, it must distinguish between "Not Recognized" (Class 0) and "Recognized" (Class 1), indicating either failure or success in facial identification [Visa et al. 2011].

The evaluation process aims to minimize classification errors and accurately categorize inputs. The confusion matrix provides a structured approach to assess prediction outcomes by classifying them into true positives, false positives, false negatives, and true negatives. From these, key performance metrics—such as accuracy, precision, recall, and F1-score—are computed, offering distinct perspectives on the model's performance. These metrics enable a comprehensive understanding of the model's strengths and limitations, identifying areas for improvement and ensuring its robustness for real-world applications. To compute these performance metrics, Scikit-learn [Pedregosa et al. 2011], a widely used Python library for ML, is utilized. Scikit-learn provides robust tools for model evaluation, including functions to generate confusion matrices and derive key classification metrics. Leveraging this library ensures a standardized and efficient assessment of the model's predictive capabilities, facilitating comparisons and optimizations where necessary.

The evaluation process utilizes 100 instances for each task, including facial images under varying lighting conditions and orientations, and fall detection data with individuals in different postures. To ensure robustness, the model is trained on open-access datasets, specifically VGGFace2 [Cao et al. 2018] for facial recognition and UR Fall Detection Dataset [Kwolek and Kepski 2014] for fall detection. As these datasets are publicly available and widely used in research, no real human participation was required, eliminating ethical concerns related to privacy and consent.

6.2. Facial Recognition

Assessing the performance of the facial recognition system is essential to understanding its reliability in distinguishing between recognized and unrecognized faces. Table 3 summarizes these results, highlighting the system's ability to correctly classify facial recognition instances. The model achieved an overall accuracy of 95%, indicating that 95 out of 100 instances were correctly classified. For the "Not Recognized" category, the model obtained a precision of 0.86 and a recall of 1.00, resulting in an F1-score of 0.92. This indicates that all actual "Not Recognized" cases were correctly classified. However, some instances predicted as "Not Recognized" were actually "Recognized", leading to a slight reduction in precision. This suggests that the model occasionally misclassifies recognized faces as "Not Recognized," causing false negatives.

Conversely, the "Recognized" category achieved perfect precision (1.00) and a recall of 0.93, yielding an F1-score of 0.96. This result indicates that all predicted "Recognized" instances were indeed correct, but 7% of actual "Recognized" cases were misclassified as "Not Recognized". The macro-average metrics—93% precision, 96% recall, and 94% F1-score—represent the simple mean across both classes, providing a balanced evaluation of the model's performance. On the other hand, the weighted average, which accounts for the class imbalance (70 instances for Class 1 vs. 30 for Class 0), yielded slightly higher values: 96% precision, 95% recall, and 95% F1-score. The increase in

Table 3. Classification metrics facial recognition

Classification report				
	Precision	Recall	F1-score	Support
0	0.86	1.00	0.92	30
1	1.00	0.93	0.96	70
Acuraccy	-	-	0.95	100.0
Macro avg	0.93	0.96	0.94	100.0
weighted avg	0.96	0.95	0.95	100.0

weighted scores suggests that the model performs better for the majority class ("Recognized"), influencing the overall evaluation. The support values (30 for Class 0 and 70 for Class 1) provide context for the dataset distribution, showing that the model was tested on a dataset where recognized faces were more frequent than unrecognized ones. This difference in class distribution is crucial for understanding the impact of class imbalance on model performance.

6.3. Fall Detection

Evaluating the effectiveness of the fall detection system is fundamental to understanding its reliability in real-world applications. Table 4 summarizes these results, demonstrating the system's ability to distinguish between fall and non-fall events. The model achieved an overall accuracy of 0.97, meaning that 97% of all predictions were correct. For individual class performance, Class 0 ("No Fall") achieved a precision of 1.00 and a recall of 0.93, leading to an F1-score of 0.96. This indicates that while all predicted "No Fall" instances were correct (precision = 1.00), some actual "No Fall" cases were misclassified as falls (recall = 0.93), suggesting a slight tendency to over-detect falls.

Table 4. Classification metrics fall detection

Classification report				
	Precision	Recall	F1-score	Support
0	1.00	0.93	0.96	40
1	0.95	1.00	0.98	60
Acuraccy	-	-	0.97	100.0
Macro avg	0.98	0.96	0.97	100.0
weighted avg	0.97	0.97	0.97	100.0

Conversely, Class 1 ("Fall") showed a precision of 0.95 and a recall of 1.00, resulting in an F1-score of 0.98. This suggests that all actual falls were correctly identified (recall = 1.00), but some non-fall instances were misclassified as falls (precision = 0.95). This trade-off between precision and recall reflects the model's bias towards prioritizing fall detection over reducing false alarms, a reasonable choice in safety-critical applications.

Additionally, the macro-average and weighted-average metrics both resulted in strong F1-scores of 0.97, reinforcing the model's balanced performance across both classes. The macro-average treats each class equally, while the weighted average accounts for the dataset distribution, where Class 1 contains more instances (60 vs. 40). The similarity between these values suggests that class imbalance had little impact on overall performance.

The support values (40 for Class 0 and 60 for Class 1) provide insights into the dataset composition. With more fall cases than non-fall cases, the model was trained on a dataset that reflects real-world scenarios where falls are more frequent in the target population.

7. Conclusion

This paper presented the development and implementation of a patient monitoring system for healthcare centers, integrating IoT technology, computer vision, and machine learning (ML). The system is designed to recognize patients through facial recognition and detect movement patterns associated with falls, utilizing computer vision-equipped cameras to send real-time alerts to caregivers and healthcare professionals.

The experimental validation confirmed the system's reliability, demonstrating its feasibility and effectiveness in practical scenarios. The results were highly promising, with the system achieving 95% precision in facial recognition and 97% accuracy in fall detection. These performance metrics indicate that the system can accurately distinguish between individuals and effectively identify fall incidents, minimizing false alarms while ensuring timely interventions. Beyond its technical performance, the proposed system exhibits versatility, making it suitable for both institutional healthcare environments and home-based monitoring. Its capability to deliver accurate predictions and enable continuous daily follow-ups enhances its practicality, particularly for elderly patients or individuals with mobility impairments. By automating patient monitoring and providing real-time alerts, the system contributes to improving patient safety, reducing caregiver workload, and fostering proactive healthcare interventions.

These findings highlight the potential of integrating AI-driven solutions in healthcare monitoring, paving the way for future advancements in smart healthcare systems. Although the system is evaluated entirely using open datasets (avoiding ethical concerns related to sensitive patient data) this approach may not fully reflect the conditions and complexities encountered in real-world clinical settings. Therefore, future research should consider validating the approach in actual healthcare environments or through realistic simulations to better assess its practical applicability. Additionally, ongoing efforts could focus on enhancing model robustness, integrating wearable sensor data, and adapting the system to a wider range of healthcare scenarios.

References

Altameem, T. and Altameem, A. (2020). Facial expression recognition using human machine interaction and multi-modal visualization analysis for healthcare applications. *Image and Vision Computing*, 103:104044.

Askar, N. A., Habbal, A., Mohammed, A. H., Sajat, M. S., Yusupov, Z., and Kodirov, D. (2022). Architecture, protocols, and applications of the internet of medical things (iomt). *Journal of Communications*, 17(11).

Bharadwaj, H. K., Agarwal, A., Chamola, V., Lakkaniga, N. R., Hassija, V., Guizani, M., and Sikdar, B. (2021). A review on the role of machine learning in enabling iot based healthcare applications. *IEEE Access*, 9:38859–38890.

Bisogni, C., Castiglione, A., Hossain, S., Narducci, F., and Umer, S. (2022). Impact of deep learning approaches on facial expression recognition in healthcare industries. *IEEE Transactions on Industrial Informatics*, 18(8):5619–5627.

Cao, Q., Shen, L., Xie, W., Parkhi, O. M., and Zisserman, A. (2018). Vggface2: A dataset for recognising faces across pose and age. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 67–74. IEEE.

Cisco, T. and Internet, A. (2020). Cisco: 2020 ciso benchmark report. *Comput. Fraud Secur.*, 2020(3):4–4.

Doan, T.-N. (2022). An efficient patient activity recognition using lstm network and high-fidelity body pose tracking. *International Journal of Advanced Computer Science and Applications*, 13(8).

Erickson, B. J. and Kitamura, F. (2021). Magician's corner: 9. performance metrics for machine learning models.

Hireche, R., Mansouri, H., and Pathan, A.-S. K. (2022). Security and privacy management in internet of medical things (iomt): A synthesis. *Journal of Cybersecurity and Privacy*, 2(3):640–661.

Koçak, H. and Çetin, G. (2021). A deep learning-based iot implementation for detection of patients' falls in hospitals. In *Trends in Data Engineering Methods for Intelligent Systems: Proceedings of the International Conference on Artificial Intelligence and Applied Mathematics in Engineering (ICAIAME 2020)*, pages 465–483. Springer.

Kwolek, B. and Kepski, M. (2014). Human fall detection on embedded platform using depth maps and wireless accelerometer. *Computer methods and programs in biomedicine*, 117(3):489–501.

Masud, M., Muhammad, G., Alhumyani, H., Alshamrani, S. S., Cheikhrouhou, O., Ibrahim, S., and Hossain, M. S. (2020). Deep learning-based intelligent face recognition in iot-cloud environment. *Computer Communications*, 152:215–222.

Mitchell, T. M. and Mitchell, T. M. (1997). *Machine learning*, volume 1. McGraw-hill New York.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.

Rajan Jeyaraj, P. and Nadar, E. R. S. (2022). Smart-monitor: Patient monitoring system for iot-based healthcare system using deep learning. *IETE Journal of Research*, 68(2):1435–1442.

Talaat, F. M. (2023). Real-time facial emotion recognition system among children with autism based on deep learning and iot. *Neural Computing and Applications*, 35(17):12717–12728.

Verma, D., Singh, K. R., Yadav, A. K., Nayak, V., Singh, J., Solanki, P. R., and Singh, R. P. (2022). Internet of things (iot) in nano-integrated wearable biosensor devices for healthcare applications. *Biosensors and Bioelectronics: X*, 11:100153.

Visa, S., Ramsay, B., Ralescu, A. L., and Van Der Knaap, E. (2011). Confusion matrix-based feature selection. *Maics*, 710(1):120–127.

Vishnu, S., Ramson, S. J., and Jegan, R. (2020). Internet of medical things (iomt)-an overview. In *2020 5th international conference on devices, circuits and systems (ICDCS)*, pages 101–104. IEEE.